Acquisition of Place Knowledge Through Case-Based Learning

PAT LANGLEY
Institute for the Study of Learning and Expertise
2164 Staunton Court, Palo Alto, CA 94306

KARL PFLEGER
Knowledge Systems Laboratory, Computer Science Dept.
Stanford University, Stanford, CA 94305

ISLE

Institute for the Study of
Learning and Expertise

Technical Report 95-2

March 15, 1995

Acquisition of Place Knowledge Through Case-Based Learning

Pat Langley◇ (LANGLEY@CS.STANFORD.EDU)
Institute for the Study of Learning and Expertise
2164 Staunton Court, Palo Alto, CA 94306

Karl Pfleger (KPFLEGER@HPP.STANFORD.EDU)
Knowledge Systems Laboratory, Computer Science Dept.
Stanford University, Stanford, CA 94305

Abstract

In this paper we define the task of place learning and describe one approach to this problem. The framework represents distinct places as evidence grids, a probabilistic description of occupancy. Place recognition relies on case-based classification, augmented by a registration process to correct for translations. The learning mechanism is also similar to that in case-based systems, involving the simple storage of inferred evidence grids. Experimental studies with physical and simulated robots suggest that this approach improves place recognition with experience, that it can handle significant sensor noise, that it benefits from improved quality in stored cases, and that it scales well to environments with many distinct places. Previous researchers have studied evidence grids and place learning, but they have not combined these two powerful concepts, nor have they used the experimental methods of machine learning to evaluate their methods' abilities.

◇ Also affiliated with the Robotics Laboratory, Computer Science Department, Stanford University, Stanford, CA 94305

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	March 15, 1995	Annual Report 3/1/94–2/28/95

4. TITLE AND SUBTITLE

Acquisition of Place Knowledge Through Case-Based Learning

5. FUNDING NUMBERS

N00014-94-1-0505 (G)

6. AUTHOR(S)

Pat Langley
Karl Pfleger

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Institute for the Study of
Learning and Expertise
2164 Staunton Court
Palo Alto, CA 94306

8. PERFORMING ORGANIZATION REPORT NUMBER

ISLE Technical Report 95-2

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Office of Naval Research
Intelligent Systems Program
800 North Quincy Street
Arlington, Virginia 22217

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

A shorter version of this paper will appear in *Proceedings of the Twelfth International Conference on Machine Learning* (1995). Morgan Kaufmann.

12a. DISTRIBUTION AVAILABILITY STATEMENT

Approved for public release; distribution unlimited

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

In this paper we define the task of place learning and describe one approach to this problem. The framework represents distinct places as evidence grids, a probabilistic description of occupancy. Place recognition relies on case-based classification, augmented by a registration process to correct for translations. The learning mechanism is also similar to that in case-based systems, involving the simple storage of inferred evidence grids. Experimental studies with physical and simulated robots suggest that this approach improves place recognition with experience, that it can handle significant sensor noise, that it benefits from improved quality in stored cases, and that it scales well to environments with many distinct places. Previous researchers have studied evidence grids and place learning, but they have not combined these two powerful concepts, nor have they used the experimental methods of machine learning to evaluate their methods' abilities.

14. SUBJECT TERMS

machine learning, case-based reasoning, probabilistic evidence grids
robotic localization, place recognition, experimental comparison

15. NUMBER OF PAGES 17

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	UL

1. Introduction and Basic Concepts

A physical agent exists in an environment, and knowledge about that environment can aid its achievement of goals. One important type of environmental knowledge concerns the spatial arrangement of the agent's surroundings. For this reason, research on the representation, use, and acquisition of spatial knowledge has occupied an important role in the field of robotics. However, work on this topic has seldom made contact with the concepts or methodology of machine learning. In this paper, we consider a novel approach to this area that incorporates ideas from both of these disciplines.

Let us begin with some definitions of concepts and tasks that appear central to spatial reasoning. Consider a physical agent, say a robot, that is situated in the world. We can say that:

Definition 1 *The* POSITION *of an agent is a coordinate in 2D or 3D space.*

Position corresponds to ground truth, giving the actual location of the agent in some established coordinate system. We might also define the related concept of agent *orientation*, but here we will assume the agent has a 360 degree field of view, making this notion unnecessary.

A physical agent does not typically have direct access to knowledge of its position, but it does have indirect information.

Definition 2 *A* SENSOR READING *is a description of the environment around the agent's position that has been filtered through its sensors.*

The information in sensor readings may be imperfect in various ways. For example, it may be incomplete in that it describes only certain characteristics of the local environment, and it may be noisy in that sensor readings for the same position may produce different results at different times.

Nevertheless, the agent must find some way to use this information to make useful inferences. This suggests a natural task for a physical agent:

Definition 3 LOCALIZATION *involves determining the position of the agent in the environment from a set of sensor readings.*

Other tasks, such as navigating from position A to position B, are certainly possible. But note that an agent cannot begin to carry out such a task without first knowing A and without knowing when it has achieved B. Thus, localization seems more basic than navigation, and we will focus our attention on it here.

However, in many situations humans seem to care less about their exact position in space than about more abstract spatial regions. This suggests another, somewhat different, concept:

Definition 4 *A* PLACE *is a contiguous sets of positions in 2D or 3D space.*

Robotics researchers have paid relatively little attention to the notion of place, but its central role in human spatial reasoning suggests it deserves a closer look. Naturally, this new concept lets us define an associated performance task by analogy with the localization task:

Definition 5 PLACE RECOGNITION *involves determining the place in which the agent currently resides from a set of sensor readings.*

At least in principle, the place recognition task seems more tractable than localization, in that it transforms a problem of numeric prediction into one involving discrete classification. One can also carry out localization within the context of a given place, but this in turn may be easier than global localization. Navigation between two places may also be simpler than navigation between two positions.

Of course, reliance on places rather than positions also introduces a problem: one must specify some descriptions in memory that let the agent map sensor readings onto place names. One might attempt to enter such descriptions manually, but it seems desirable to automate this process, suggesting a final task:

Definition 6 PLACE LEARNING *involves the induction of descriptions, from the sensor readings and place names for a set of training positions, that let the agent accurately recognize the places of novel positions.*

Note that this task formulation makes minimal demands on the teacher, who does not have to give the agent information about its actual positions. Rather, the agent collects its own sensor readings, and the teacher must only label each reading as an instance of one place or another. This formulation assumes supervised training data, but unsupervised versions, in which the agent decides on its own place names, are possible as well. We will touch briefly on unsupervised place learning in Section 5, but we will focus on the supervised version in this paper.

In the pages that follow, we present one approach to dealing with knowledge about places. First we describe a representational formalism for storing place knowledge – evidence grids – and then examine a method for place recognition that operates on this representation. After this, we consider a simple learning process that lets one acquire and refine knowledge of places. Next we present some hypotheses about our approach, along with some experimental tests of those hypotheses. Finally, we review related work on spatial learning and discuss some directions for future work.

One important difference between our approach and earlier robotics work on spatial knowledge lies in our incorporation of ideas from machine learning. In particular, we view place recognition as a classification task and we view place learning as a supervised concept induction task. This suggests not only certain learning methods, but also the use of experimental methods prevalent in machine learning to evaluate our technique. However, the tasks of place recognition and place learning introduce some difficulties not usually present in machine learning research, such as the pervasive presence of significant sensor noise. Our approach to the problem is designed with these issues in mind.

2. Representation, Use, and Acquisition of Place Knowledge

With the above definitions in hand, we can examine one approach to learning place knowledge. However, before we address the acquisition process, we should first consider the manner in which we represent knowledge about places and the performance element that takes advantage of that spatial knowledge base.

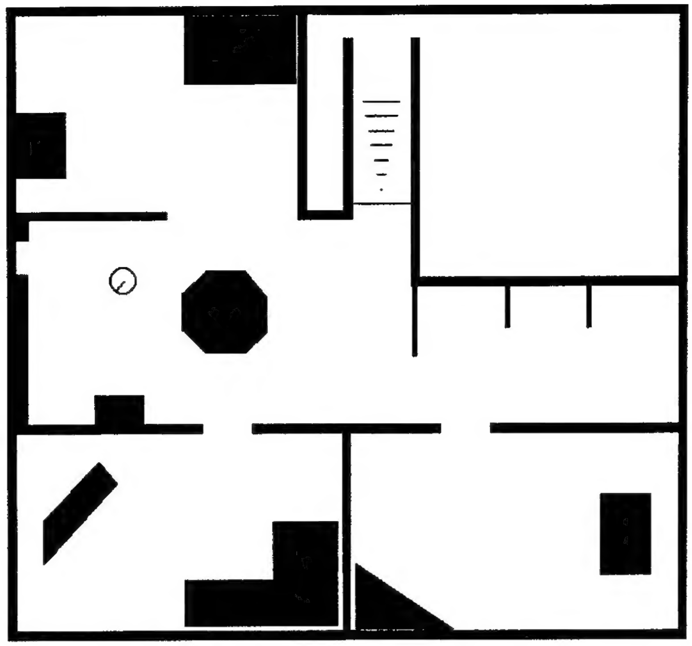

Figure 1. A simulated office environment with a number of distinct places.

2.1 The Evidence Grid Representation

Robotics researchers have explored a variety of formalisms for representing spatial knowledge. One approach relies on geometric primitives to describe the edges or surfaces of obstacles in the environment. For example, one can use a set of lines to approximate the walls of an office and the furniture it contains. Such representations are precise, but Schiele and Crowley (1994) note that they can be difficult to use when sensors are noisy.

Another common scheme involves dividing the environment into a rectangular grid of mutually exclusive cells, each corresponding to a distinct position in space. In this framework, each cell is specified as either occupied (containing an obstacle) or open (containing none). This approach is well suited to navigation tasks in which one already knows the structure of the environment (i.e., which cells are occupied) and the position of the agent within the grid. However, this scheme is not designed to handle the uncertainty that arises when the position is unknown or when the agent has yet to learn the structure of the environment.

An alternative framework uses the *evidence grid* (Elfes, 1989; Moravec & Blackwell, 1992), a data structure that is specifically designed to deal with uncertainty. In this approach, each cell C has an associated probability that C is occupied by some tangible object that would block the agent's path if it tried to move through the cell. These probabilities range from near zero (nearly certain a cell is open) to near one (nearly certain a cell is occupied), with the middle corresponding to cells for which little information is available (e.g., behind a wall or inside an object). We will adopt this framework in the current paper.

Figure 1 shows the position of an agent in a room within a larger office environment, similar in structure to an actual area at Stanford University. Figure 2 depicts evidence grids constructed from simulated sensor readings taken from positions (a) in the top left room and (b) in the lower

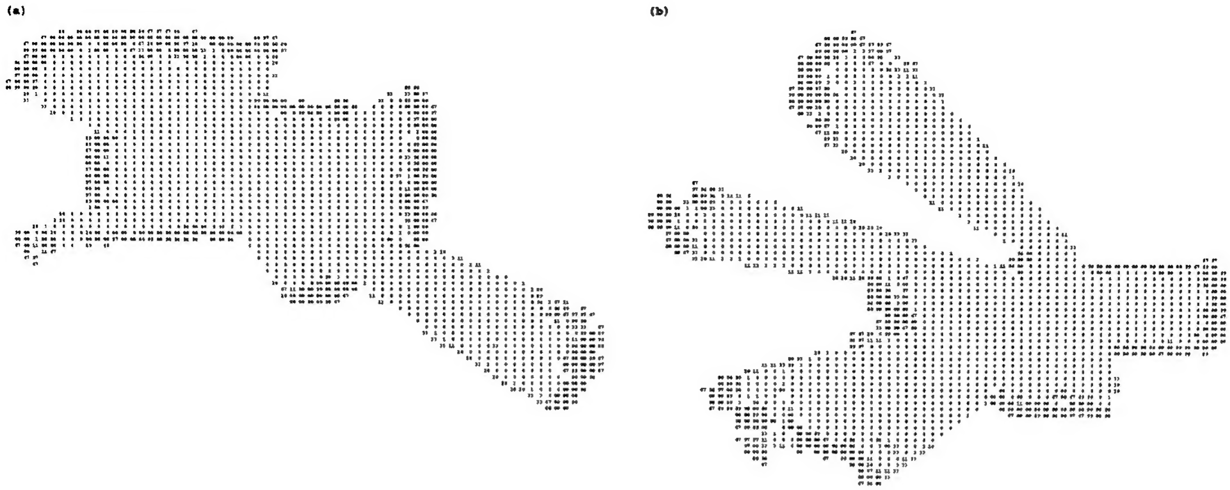

Figure 2. Evidence grids generated from simulated sonar readings for (a) the top left room in Figure 1 and (b) the lower left room in the same figure.

left room from the same orientation. Note that open regions within the agent's view have low probability of being occupied (lighter shades) and that edges of obstacles and walls within view have high probability (darker shades). However, areas that are occluded, such as those behind obstacles and walls, have probabilities around $\frac{1}{2}$ (empty regions), since the agent's sensors provide no information about them. Of course, the agent can construct a more complete evidence grid by moving around the environment to collect sensor readings from different viewpoints.

Previous work with evidence grids has emphasized their use in representing single rooms over a relatively short period. However, they also have potential for handling large-scale spatial knowledge over longer time spans. An agent could store its knowledge about an entire building or even a city in a single, large evidence grid. But this scheme seems impractical due to the difficulties inherent in integrating information from distant regions into a single map.

A more tractable approach to representing large-scale spatial knowledge, which we take here, involves storage of separate evidence grids for each distinct place. For example, one might use a different grid to encode each room in a building. This knowledge can be augmented by geometric relations among places, which would support navigation planning, but we will not address that aspect here. The retention of place descriptions in memory has much in common with the storage of a *case library* in work on case-based reasoning. In both frameworks, the stored items represent alternative situations in which the agent can find itself, and which suggest different inferences.

2.2 Case-Based Recognition of Places

Now that we have described the nature of evidence grids, we can examine their use in place recognition. Let us assume the agent has a stored place library, with each place described as an $R \times C$ evidence grid with an associated place name. Our approach to place recognition relies on a three-step process that has much in common with case-based reasoning.

First, the agent constructs a temporary or short-term evidence grid for its current position from a set of sensor readings. This involves transforming each sensor reading into a probability of

occupancy for each cell. Following Moravec and Blackwell (1992), we assume a sensor model that specifies this mapping. The result is a temporary evidence grid, based on the sensor reading, that characterizes the region in the vicinity of the agent. The agent may repeat this process a number of times, in each case incorporating the result into the temporary grid using a Bayesian updating scheme. We will not describe this updating process in depth, but readers can find details in Elfes (1989) and in Moravec and Blackwell (1992).

Next, the agent matches the short-term evidence grid against each of the grids stored in the place library. The evaluation function used in this comparison process measures the degree of match between two grids. Specifically, if $S_{r,c}$ is the probability associated with the rth of R rows and the cth of C columns for the short-term grid, and if $L_{r,c}$ is the analogous probability for the stored, long-term grid, then

$$M = \sum_{r}^{R} \sum_{c}^{C} F(S_{r,c}, L_{r,c})$$

computes the similarity between the short-term and stored grids. One can instantiate the function F in many ways, provided they satisfy certain properties: two cells should be treated as similar if they are confident in the same direction, as dissimilar if they are confident in opposite directions, and generally ignored if either is uncertain.

Moravec and Blackwell (1992) implement this cell-to-cell component of grid similarity as the expression

$$F(S_{r,c}, L_{r,c}) = log\,[S_{r,c}L_{r,c} + (1 - S_{r,c})(1 - L_{r,c})] + 1 \quad ,$$

which varies from one (a perfect match) to negative infinity (the worst possible match). Reflection suggests that this scheme might give very low match scores to reasonably similar grids if even a few cells are confident in opposite directions. For this reason, we decided to use an alternative metric:

$$F(S_{r,c}, L_{r,c}) = \begin{cases} 1 & \text{if } S_{r,c} > \frac{2}{3} \text{ and } L_{r,c} > \frac{2}{3} \\ 1 & \text{if } S_{r,c} < -\frac{2}{3} \text{ and } L_{r,c} < -\frac{2}{3} \\ -1 & \text{if } S_{r,c} > \frac{2}{3} \text{ and } L_{r,c} < -\frac{2}{3} \\ -1 & \text{if } S_{r,c} < -\frac{2}{3} \text{ and } L_{r,c} > \frac{2}{3} \\ 0 & \text{otherwise .} \end{cases}$$

We felt this measure would be less sensitive to situations in which disagreements arise between cells having high certainty, thus eliminating the problem predicted for the Moravec/Blackwell measure.

The above metrics assume that the stored and temporary grids are described in the same coordinate system. One can plausibly assume the presence of a reasonably accurate compass to determine the relative rotations, but possible differences in translation requires some form of registration that coerces the temporary evidence grid into the same coordinate system as the stored place. To this end, our system carries out an exhaustive search using operators that modify the position by one grid row or column, evaluating each alternative using the metric M defined above.[1] The system selects the translation that gives the highest M score; the resulting registered grid localizes the

1. When translation causes two grids to overlap on only $R' \times C'$ cells, the metric uses only these cells in its summation. This creates a bias toward stored grids that share more cells with the temporary grid, which seems reasonable, but it does not actively punish a stored grid for having only partial overlap.

agent with respect to that grid. If the compass is not accurate, one can extend this approach to correct for small offsets in rotation.

Finally, the agent compares the match scores for the various registered grids and selects the best of these competitors. This strategy provides both the place name associated with the selected evidence grid and the estimated position within that place description. Because adjacent evidence grids may cover overlapping regions, this scheme has some potential for misclassifying a place based on its outlying rather than its central cells. An alternative strategy would let the agent associate distinct place names with different cells in the same stored grid, then predict the name specified for the cell nearest to the estimated position. However, this issue has not been a problem in our studies to date, so our current system relies on the simpler classification strategy.

As we noted earlier, this approach has much in common with methods for case-based reasoning. Here the evidence grids in the place library correspond to stored cases, whereas the short-term grid maps on to a test case for which one wants to make a prediction. The match function corresponds to the similarity metric that determines the nearness of the test case to each stored case in an $R \times C$ dimensional space, and the final classification step is similar to that used in the nearest neighbor method, perhaps the simplest case-based technique. The fact that each evidence grid may be a probabilistic summary, computed from a set of sensor readings, differs from the prototypical case-based system, but abstract cases are not that unusual. A more intriguing difference concerns the registration process. Many case-based systems incorporate some *adaptation* method, but usually this occurs after case selection, whereas here adaptation (registration) takes place during the evaluation (match) process itself.

2.3 Case-Based Learning of Place Knowledge

Now let us consider an approach to learning place knowledge that is stored as evidence grids. We would like an incremental process, since the agent encounters its environment sequentially. However, we are not concerned here with the task of effectively exploring an unknown world, so we will assume that the agent is led to a position, given time to observe its surroundings, given a place name, led to another position, and so forth.

Given our commitment to a place library and to a case-based method for place recognition, we naturally assume a case-based learning scheme as well. In particular, at each position to which it is led, the agent constructs a short-term evidence grid S using the method described above. The system then simply adds the new grid to the place library, along with specified place name. The same place name may be associated with multiple evidence grids, but this seems appropriate if they produce different sensor readings.

At first glance, this approach to place learning sounds guaranteed to work, in that one simply stores a description for each place, after which recognition will be perfect. However, this view ignores the central feature of the task – uncertainty. Even with noise-free sensors, the same place typically looks different from different positions, if only because different regions are occluded. Moreover, standard robotic sensors such as sonar are notoriously noisy, and will produce different sensor readings, and thus different evidence grids, even when repeated from the same position.

In addition, the dimensionality of the resulting space is very high, with one attribute for each cell in the $R \times C$ evidence grid. Langley and Sage (in press) have shown that the learning rate of case-based methods can be drastically slowed by the presence of irrelevant attributes. Since typical rooms contain large open areas, it seems plausible that the cells that describe such areas will make place learning difficult. Thus, the adequacy of this approach remains an open question that is best answered by experiment.

3. Experimental Studies of Place Learning

In Section 1 we formulated the place learning task in terms similar to those used to describe other induction problems. Thus, we can use the experimental methods developed for machine learning to evaluate the robustness of our framework. In this section we present a number of hypotheses about the system's behavior, followed by experimental tests of those hypotheses. Our primary measure of performance was recognition accuracy for places in a test set of evidence grids that differ from those in the training set.

3.1 The Experimental Setting

The experiments we designed to evaluate the abilities of our approach relied on both a physical robot – a Nomad 200 with a 16-sensor sonar ring – and a high-fidelity simulation of this machine. The physical environment was a suite of offices and common areas at Stanford University, and the simulated environment was an idealized layout of a similar suite, depicted earlier in Figure 1. We used the physical Nomad to ensure realism in our results, while the simulation gave us experimental control over device parameters not possible with the actual robot.

We generated each training or test case by placing the physical or simulated robot in a position, collecting readings from the sonar ring to construct an initial evidence grid, rotating and/or moving the robot (as described below), collecting new sonar readings and updating the evidence grid, and repeating this process many times. For the simulated robot, we generated six different grids for each of six distinct places,[2] giving 36 total evidence grids. For the physical robot, we produced only three grids for each place (because the process took longer), giving 18 total grids.

The Nomad simulator incorporates a number parameters that affect the quality of sonar information. For example, the `error` parameter controls random variation in the distance returned by the sonar sensors, `critical` controls the angle of incidence at which specular reflection occurs, and `halfcone` controls the angular width of each sonar signal. Unless otherwise specified, we set `error` to 0.15, which was our best estimate of the error encountered by the physical robot, and we left all other parameters at their default values, which produce a 25 degree field of view for each sensor and specular reflection at angles of incidence with the sensed surface of 30 degrees or less.

For each experimental condition with the simulated environment, we ran the learning system 400 times with different random partitions of the evidence grids into 33 training and three test cases,

2. These places corresponded to the lower left, lower right, middle right, and upper left rooms in Figure 1, and to the areas to the left and right of the octagonal table in the figure.

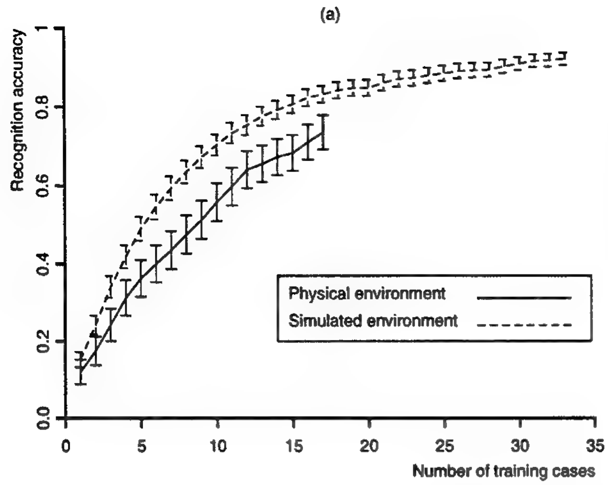

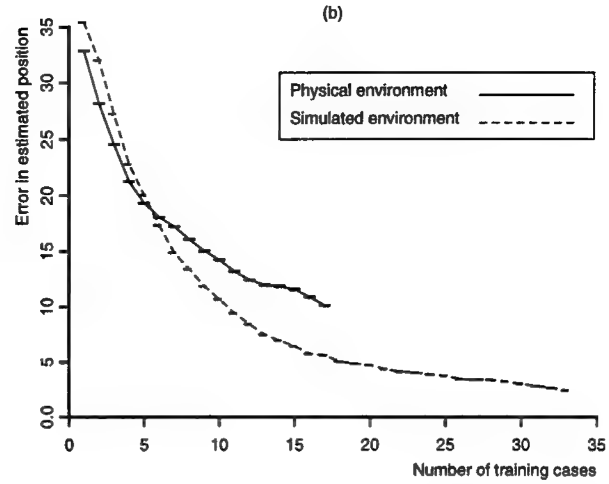

Figure 3. Learning curves with 95% confidence intervals for the case-based place learning system for a physical Nomad robot and a simulated robot in a similar office environment, (a) using recognition accuracy as the performance measure and (b) using error in estimated position.

randomly ordering the storage of training cases. For the physical environment, we partitioned the grids into 17 training cases and one test case, again averaging over 400 runs for each condition.

3.2 Improving Place Recognition with Experience

Following Kibler and Langley (1988), we can divide the factors that affect the learner's behavior into two broad types, those involving characteristics of the environment and those involving features of the learner. The most basic environmental characteristic is the number of training cases available. Naturally, we hypothesized that the accuracy of place recognition would improve as the agent encounters more positions. However, the literature sometimes reports actual decreases in performance, so we needed to explicitly test this expectation.

Figure 3 (a) shows the learning curve, giving 95% confidence intervals, for the physical Nomad robot when each training and test grid was based on 45 sonar readings, taken from a single position but with successive orientations incremented by one degree. As expected, the system's ability to recognize places gradually increases as it observes and stores more training cases. However, the shape of the curve suggests that the learning task is not trivial, in that multiple cases for each place are needed to achieve even 70% accuracy. The curve has not yet leveled off at 17 instances, so presumably additional cases would further improve recognition.

Close inspection of the inferred structures reveals that, from certain views, the registered evidence grids for two different places occasionally appear more similar than the grids for two different positions within the same place. This should not be surprising, given the noise inherent in sonar sensors and given that objects can occlude portions of a place from some positions. Table 1 shows the actual confusions that occur, on average, for the physical robot after 17 training cases. Note that most errors involve the misclassification of place (e), which is confused with places (a) and (c), and the mislabeling of (d), which is classified as (c).

Table 1. Confusion matrices with probabilities of labeling for the six places used in the experiments. Rows indicate the correct place names, whereas columns show the predicted place after training.

PHYSICAL ROBOT						
	(a)	(b)	(c)	(d)	(e)	(f)
(a)	1	0	0	0	0	0
(b)	0	1	0	0	0	0
(c)	0	0	1	0	0	0
(d)	0	0	$\frac{2}{3}$	$\frac{1}{3}$	0	0
(e)	$\frac{2}{3}$	0	$\frac{1}{3}$	0	0	0
(f)	0	0	0	0	0	1

SIMULATED ROBOT						
	(a)	(b)	(c)	(d)	(e)	(f)
(a)	$\frac{5}{6}$	0	$\frac{1}{6}$	0	0	0
(b)	0	1	0	0	0	0
(c)	0	0	1	0	0	0
(d)	0	0	$\frac{1}{6}$	$\frac{5}{6}$	0	0
(e)	0	0	$\frac{1}{6}$	0	$\frac{5}{6}$	0
(f)	0	0	0	0	0	1

Figure 3 (a) also shows an analogous learning curve for the simulated robot. The general shape of the curve is very similar to that for the physical device, but the rate of learning is somewhat higher. Table 1 also shows the averaged confusion matrix for this experimental condition after the learner has seen 35 training cases. Although a few errors still occur, the performance component generally assigns the correct place name to the test cases.

Naturally, we are interested not only in our method's ability to recognize places, but its ability to identify the precise position of the robot within a given place. Thus, we also measured the absolute difference between the actual robot position in each test case and the estimated position as computed during the registration process. Figure 3 (b) shows the learning curves for the physical robot, as well as similar results for the simulated one. In this case, since we are measuring error rather than accuracy, the quantities start high and gradually decrease with experience. Again, behavior in the simulated environment generally mimics that in the physical world, though the system fares somewhat better on the former.

3.3 Sensor Noise and Case Quality

The amount of sensor noise constitutes a more interesting environmental factor. We would not expect increased noise to affect the asymptotic accuracy, but it should decrease the rate of place learning, that is, the number of training cases needed to reach a given accuracy level. Nevertheless, we hope that the probabilistic nature of evidence grids will let our approach degrade gracefully with increasing amounts of sensor noise.

Figure 4 shows two evidence grids constructed from simulated sonar signals collected from the same position and orientation within the lower left room in Figure 1. For grid (a), we set the simulator's `error` parameter to zero, so that there was no sensor noise. For grid (b), we set this parameter to 0.45, producing significant noise. The resulting evidence grids are similar but contain noticeable differences, suggesting that the basic inference method is robust but that sensor noise also has some effect.[3]

3. The inferred "arms" in (b) appear to be artifacts of the grid updating scheme; the more generic effect of sensor noise is to create more ragged boundaries along the edges of objects.

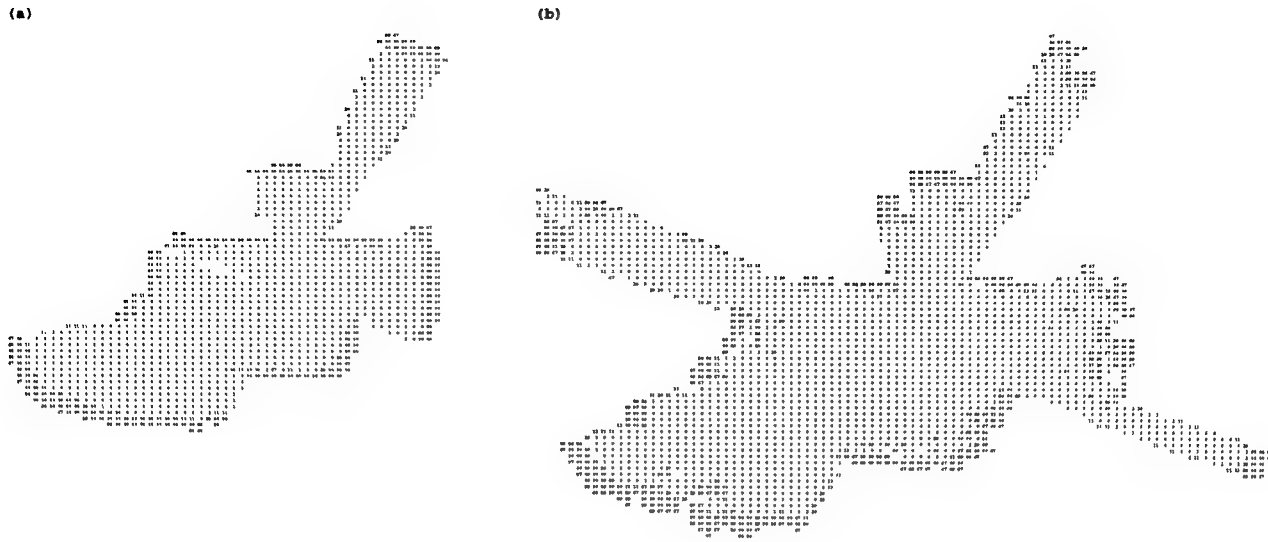

Figure 4. Evidence grids generated from 45 simulated sonar readings for the lower left room in Figure 1 using (a) zero sensor noise and (b) a 0.45 noise setting.

Fortunately, our reliance on evidence grids suggests a natural response to noisy sense data. Because each stored case can be based on multiple sensor signals, we can attempt to improve the *quality* of these cases by increasing the number, or altering the arrangement, of the signals used to generate them. We hypothesize that place descriptions based on more sensor readings will be less affected by increases in sensor noise. Thus, we predict an interaction between these two independent variables, specifically one that affects learning rate but not asymptotic accuracy.

To test this hypothesis, we used the Nomad simulator to produce four different levels of sensor noise, in which the `error` parameter was set to 0.0, 0.15, 0.30, and 0.45, respectively. We also attempted to vary the quality of the stored cases by using two different sensing strategies. In one, we based each evidence grid (both training and test cases) on 45 sonar readings collected from a single position but produced at orientations one degree apart, as used to generate the results in Figure 3. In the other, we based each grid on 90 readings, produced by repeating this strategy in two nearby positions within the same room.

Figure 5 (a) shows the learning curves that result for the zero and 0.45 noise levels using the one-position sensing strategy, whereas Figure 5 (b) presents analogous results for the two-position strategy. (The results for the 0.15 and 0.3 settings fell between these extremes; we have omitted them for the sake of clarity.) The two-position scheme clearly fares better than the simpler strategy, but the curves diverge somewhat from our predictions. The rate of learning for the two-position method is much higher than for the one-position method, even when no sensor noise is present. Also, the introduction of sensor noise clearly affects both strategies, but it alters only the learning rate for the more sophisticated scheme, while it actually appears to reduce the asymptotic recognition accuracy for the simpler one.

The general superiority of the two-position strategy is hardly surprising, in that its evidence grids are based on twice as many sonar readings. Ideally, we would prefer a sensing scheme that is robust with respect to noise but that requires no more sensing than the initial strategy. To

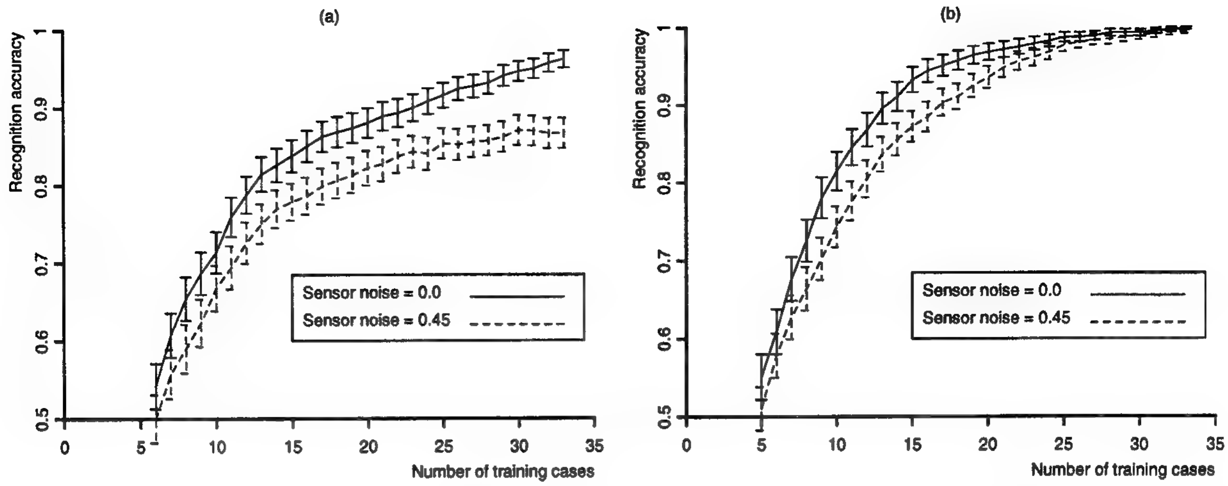

Figure 5. Learning curves for the case-based place learning system for two levels of sensor noise when evidence grids are based on (a) 45 readings from one position and (b) 90 readings from two nearby positions.

this end, we explored a third method that takes three sonar readings at an initial position, with rotational increments of 7.5 degrees, moves a fixed amount and takes another three readings in the same manner, then repeats this process until completing a total of 45 readings. The resulting evidence grid is based on sensing over the entire straight-line path, reducing the chance of occlusion and hopefully reducing the effect of sensor noise.

Figure 6 (a) presents learning curves for this sensing strategy on the two of the simulated noise levels. For the noise-free situation, the behavior is nearly identical to that for the 90-reading strategy, even though grids are based on half as many sonar signals. However, sensor noise significantly degrades this strategy's behavior, though its accuracies remain well above those for the one-position method. Clearly, basing evidence grids on a number of distinct positions within a given place gives better results than basing them on one position, but increasing the number of readings also has desirable effects. It seems likely that more sophisticated sensing strategies, which sample readings in a more intelligent manner, would produce even better results.

3.4 Effect of the Similarity Metric

In Section 2.2 we described the similarity metric used to assign a short-term evidence grid to the stored grid that best matches it. This metric sums over the cells on which the two grids overlap, using a function F to measure the similarity of individual cells. We contrasted our implementation of F, which takes on the values 1, 0, and -1, with the implementation used by Moravec and Blackwell (1992), which ranges from one to negative infinity. We presented some intuitive arguments for preferring our formulation, but the question of which measure behaves better in practice is ultimately an empirical one.

Figure 6 (b) presents experimental results for the two similarity metrics, using training and test cases from six places based on 45 simulated sonar readings from one position.[4] The learning curve

4. The setting for `halfcone` here was 25 degrees, rather than 12.5, as in other runs. The `error` was again set to 0.15, and the other parameters remained unchanged.

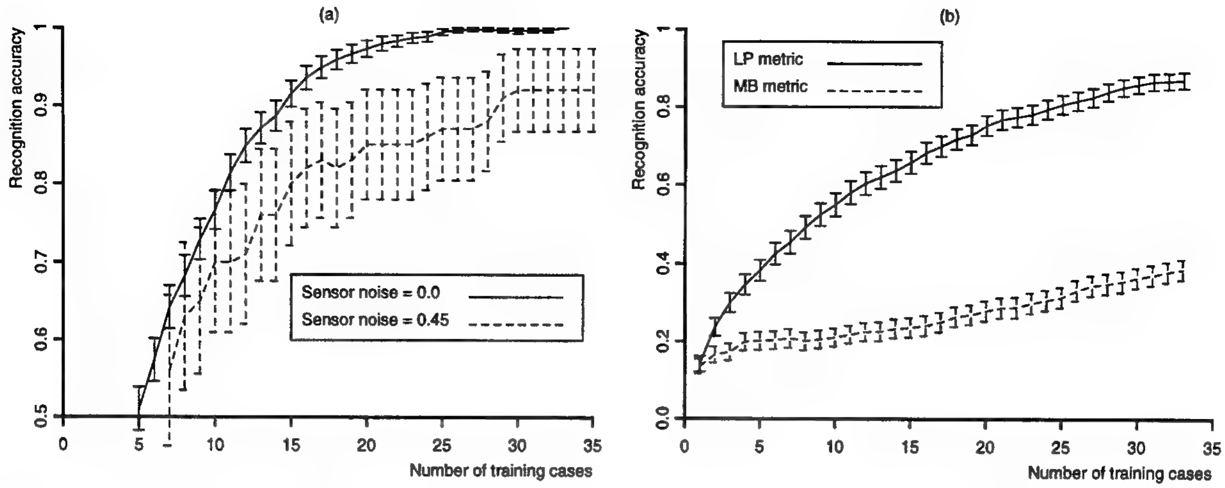

Figure 6. (a) Learning curves for two levels of sensor noise when evidence grids are based on 45 readings taken at equal intervals along a line between two positions. (b) Learning curves using the Langley/Pfleger (LP) similarity metric and the Moravec/Blackwell (MB) metric.

for our version of the F function is similar to those we have seen earlier in the paper. In contrast, the curve for the Moravec and Blackwell metric reveals learning at a much slower rate, reaching only 39% accuracy after 33 training cases, as compared with 87% for our measure. These results do not imply that our approach is the only viable option, but they do show that the similarity measure can make a substantial difference in place recognition, and that our metric performs much better than one proposed alternative.

3.5 Number of Distinct Places

Some real-world environments contain many distinct places, making it desirable for a learning method to scale well as the number of places increases. We obtained preliminary results along these lines by examining our algorithm's behavior with different subsets of the places available in our environment. Figure 7 (a) shows the learning curves that result for two through six places, with each case based on 45 simulated sonar readings from one position. Each reported accuracy is averaged over 400 runs for each possible subset of k out of six places, using 35 randomly selected training cases and one test case. Thus, when $k = 2$ we carried out $\binom{6}{2} \times 400 = 6000$ runs, and when $k = 3$ we carried out $\binom{6}{3} \times 400 = 8000$ runs. We have not reported confidence intervals here, since the accuracies are averages of averages.

Naturally, increasing the number of places decreases the speed of learning, but we can also examine the rate of this decrease. Note that the figure also shows where each learning curve crosses the level of 90 percent accuracy. These crossover points produce the *scaling* curve in Figure 7 (b), which maps the number of distinct places against the number of training cases needed to reach this accuracy level. This higher-order curve seems to be either linear or quadratic, but the analogous scaling curves for the two-position and straight-line sensing strategies, also shown, definitely appear linear. These results suggest that our approach requires, more or less, a fixed number of training cases per place, independent of the total number of places. This encourages us to believe that the method will scale well to domains that involve many more different places than the six we have examined, though ultimately we should test this prediction using larger environments.

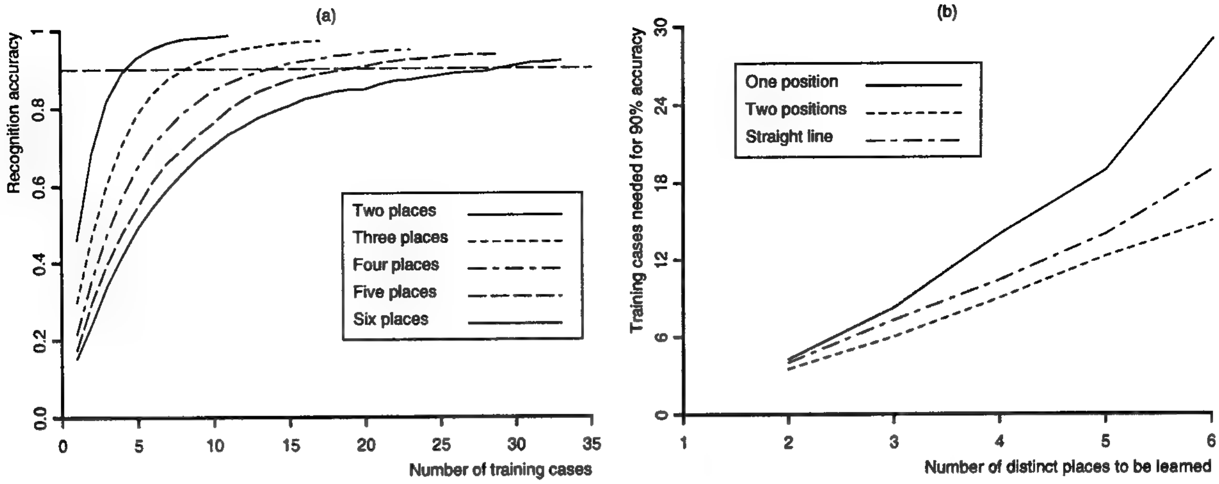

Figure 7. (a) Learning curves for different numbers of distinct places, based on 45 sonar readings from one position. (b) Scaling curves that map, for different sensing strategies, the number of training cases needed to achieve 90% accuracy as a function of the number of places.

3.6 Summary of Experimental Results

In this section, we reported on a number experiments designed to evaluate our case-based approach to the acquisition of knowledge about places. We adopted a method common in research on machine learning, stating explicit hypotheses and running experiments designed to test them. In each case, we varied one or two independent variables and observed their effects on some performance measure.

The experiments revealed a number of encouraging behaviors. Our approach to place learning generally improves its recognition accuracy as it observes more training cases, with similar results occurring for both the physical and simulated robot. The learning rate slows in the presence of sensor noise, but one can mitigate this effect by increasing the quality of the inferred evidence grids. The rate of learning also slows with increasing numbers of places, but no more than expected in any multiclass learning situation. In addition, we found that our similarity metric performs significantly better than another metric proposed in the literature.

Clearly, there exist many other factors that could influence the behavior of our place-learning method. These include the resolution of the evidence grids and the distinctiveness of the places one must learn to distinguish. However, we will reserve these issues for future studies, as the current experiments have been sufficient to show that our approach is a promising one.

4. Related Work on Learning Spatial Knowledge

Our research on the acquisition of spatial knowledge is certainly not the first in this area. Clearly, our work owes a strong intellectual debt to Elfes (1989), Moravec and Blackwell (1992), and other developers of the evidence grid framework. Our basic representation and our performance system directly employ techniques developed by these researchers. However, most research in this framework has focused on the construction of a single global map, rather than a collection of evidence grids for distinct places. Although such approaches clearly acquire spatial knowledge, they do not

involve induction in the sense of using training instances to improve performance on novel test cases, whereas our work on place learning fits easily into this paradigm. Thrun (1993) has used reinforcement learning to improve sensor interpretation for evidence-grid construction, but his goal was to construct a global map. Mahadevan (1992) describes a method that forms generalizations expressed as evidence grids, but his aim was to learn not places but action models.

Nevertheless, some researchers outside the evidence grid formalism have studied place learning. For example, Yamauchi and Beer (1994) describe ELDEN, a system that represents places in terms of means and variances of direct sensor readings, rather than inferred grid occupancies. Their place descriptions also include features for the robot's position as estimated through dead reckoning and connections to recently visited places. Place recognition involves passing each attribute's value through Gaussian functions associated with each place, then selecting the competitor with the highest sum. Learning consists of updating the means and variances for recognized places, creating new places when no existing ones match well enough, and adding predictive connections between places. Yamauchi and Beer's reliance on a Gaussian distance metric makes their method similar to our case-based approach, though ELDEN differs in its use of instance averaging, its use of raw sensor data, and the unsupervised nature of the learning processes.[5]

Lin, Hanson, and Judd (1994) have taken a similar approach to representing and using spatial knowledge. Their system also describes places (which they call *landmarks*) as means and variances of sonar readings and uses a Gaussian metric to determine the degree of match against the current sensor signals. However, their learning mechanisms include not only the creation and updating of place descriptions, but also a reinforcement process designed to improve estimates of the robot's location. This latter technique can lead the learner to add a new place or remove an existing one if these actions reduce errors in location estimates.

Kuipers and Byun's (1988) NX system also operates on direct sensory readings, but it stores only places that are distinctive in terms of optimizing certain measures. For example, NX defines the central point in a hallway corner as being symmetrical and being equidistant from the walls, in addition to containing information about the angles and distances to obstacles. The system also describes edges, which connect distinctive places, in terms of length, width, and similar characteristics. Whenever NX encounters a local optimum L on one of its measures, it compares the sensor readings to each known place P stored in memory; if the descriptions for L and P are similar, and if their locations are metrically or topologically close, the system classifies L as the known place P. Otherwise, NX creates a new place based on L's description and stores this in memory, along with its edge connections to other places. Mataric (1991) describes a similar scheme, though the details of place creation are different.

In methodological terms, Kortencamp and Weymouth's (1994) work is perhaps the most similar to our own. Their approach emphasizes *gateways* such as doors that connect two regions, but their system represents these locations using a grid structure and they evaluate its behavior in terms of recognition accuracy. However, their scheme uses hand-coded descriptions for a few gateway types

5. Yamauchi and Beer's system also incorporates an evidence grid representation, but it constructs a global map and uses this map for correcting errors in dead reckoning rather than for place recognition.

to recognize candidate places and create new ones, rather than actual supervised training data, and they compare a number of different recognition strategies, including one that combines evidence from sonar and visual sensors.

On another dimension, our approach is most similar to Yeap's (1988) work on spatial reasoning. His framework also posits the storage of distinct places, the descriptions of which are not direct sensory readings but inferred summaries. However, his "absolute space representation" does not take the form of evidence grids but rather consists of a connected sequence of line segments that, except for occasional openings, enclose an area. Yeap does not describe a performance element that uses these descriptions in place recognition but, as in our own framework, learning involves the simple storage of the inferred place descriptions, which suggests the use of a case-based method.

5. Concluding Remarks

Although our experimental studies of place learning have revealed some insight into our approach, clearly more work remains to be done. The most immediate extension would replace the current supervised learning method with an unsupervised one. Such a system must identify distinctive places on its own, as it cannot rely on a tutor for this information. To this end, we plan to employ a technique similar to that used by Yamauchi and Beer (1994), but adapted to operate on evidence grids rather than direct sensor descriptions. As the agent moves through the environment, it would regularly stop and construct a short-term evidence grid, merging this with the previous place description if the match is high enough and using the short-term grid as the basis for a new place otherwise. Discontinuities caused by passage through doors and past obstacles should be enough to identify distinguishable places.

Most methods for place learning, including those discussed above, also construct topological maps that connect different places. Clearly, this is another important direction in which to extend our approach. We expect that storing rough estimates of the direction of movement between one place and its successor will be sufficient for many navigation tasks. Upon executing a navigation plan, the agent would still need to register its location upon entering each place along the path, but expectations about the next place and its rough translation should greatly simplify the registration process. Storing recently visited places with each evidence grid could also reduce recognition errors in domains with perceptually similar places.

In future work, we also hope to develop methods for detecting distinctive features in evidence grids that would simplify the place recognition process. We envision such features as being configurations of grid cells with large differences in their probabilities, such as might occur along a wall or at a door. The recognition mechanism would use the presence of these features as cues during retrieval of candidate places and during registration, and the learning process would use the features to index places in memory. Such learned features could also play the role of landmarks, in the sense used by Levitt, Lawton, Chelberg, and Nelson (1987), that qualitatively distinguish places. One simple approach to detecting useful configurations of grid cells would draw on recent methods for feature selection in case-based learning (e.g., Langley & Sage, 1994), which use estimates of accuracy obtained through cross validation to direct search through the space of feature combinations.

In summary, we have presented a framework for representing, using, and learning knowledge about places in which evidence grids play a central role. Our approach draws on earlier work for updating these probabilistic summaries, but diverges from previous schemes by storing a set of local grids in a case library, then retrieving and matching them for use in place recognition. Experimental studies adapted from the machine learning literature indicate that this approach improves recognition accuracy with experience, that sensor noise degrades the learning process, and that improving the quality of stored cases can offset this effect. The experiments also revealed that our method scales well to increased numbers of places, and that some of its power comes from the particular similarity metric used in the matching process. Many other environmental and system factors remain to be examined, but the basic approach appears promising and suggests many natural extensions.

Acknowledgements

We owe thanks to Hans Moravec for making his evidence grid software available for use, and to Brian Yamauchi for integrating that code with the Nomad interface. Alan Schultz and Brian Yamauchi made many useful suggestions, and Barbara Hayes-Roth and Nils Nilsson provided access to a Nomadics robot. This research was supported in part by Grant Number N00014-94-1-0505 from the Office of Naval Research to the Institute for the Study of Learning and Expertise.

References

Elfes, A. (1989). Using occupancy grids for mobile robot perception and navigation. *IEEE Computer Magazine*, June, 46–58.

Kibler, D., & Langley, P. (1988). Machine learning as an experimental science. *Proceedings of the Third European Working Session on Learning* (pp. 81–92). Glasgow: Pittman.

Kortencamp, D., & Weymouth, T. (1994). Topological mapping for mobile robots using a combination of sonar and vision sensing. *Proceedings of the Twelfth National Conference on Artificial Intelligence* (pp. 979–984). Seattle, WA: AAAI Press.

Kuipers, B., & Byun, Y. T. (1988). *Proceedings of the Eighth National Conference on Artificial Intelligence* (pp. 774–779). St. Paul, MN.

Langley, P., & Sage, S. (1994). Oblivious decision trees and abstract cases. *Working Notes of the AAAI94 Workshop on Case-Based Reasoning* (pp. 113–117). Seattle, WA: AAAI Press.

Langley, P., & Sage, S. (in press). Scaling to domains with many irrelevant features. In R. Greiner (Ed.), *Computational learning theory and natural learning systems* (Vol. 4). Cambridge, MA: MIT Press.

Levitt, T. S, Lawton, D. T., Chelberg, D. M., & Nelson, P. C. (1987). Qualitative landmark-based path planning and following. *Proceedings of the Sixth National Conference on Artificial Intelligence* (pp. 689–694). Seattle, WA: AAAI Press.

Lin, L., Hanson, S. J., & Judd, J. S. (1994). *On-line learning for landmark-based navigation* (Technical Report No. SCR-94-TR-472). Princeton, NJ: Siemens Corporate Research, Learning Systems Department.

Mahadevan, S. (1992). Enhancing transfer in reinforcement learning by building stochastic models of robot actions. *Proceedings of the Ninth International Conference on Machine Learning* (pp. 290–299). Aberdeen: Morgan Kaufmann.

Mataric, M. J. (1991). Behavioral synergy without explicit integration. *Sigart Bulletin*, *2*, 130–133.

Moravec, H., & Blackwell, M. (1992). Learning sensor models for evidence grids. *Robotics Institute Research Review*. Pittsburgh, PA: Carnegie Mellon University.

Schiele, B., & Crowley, J. L. (1994). A comparison of position estimation techniques using occupancy grids. *Robotics and Autonomous Systems*, *12*, 163–171.

Thrun, S. B. (1993). Exploration and model building in mobile robot domains. *Proceedings of the IEEE International Conference on Neural Networks*. San Francisco: IEEE.

Yamauchi, B., & Beer, R. (1994). *Exploration and spatial learning in dynamic real-world environments* (Technical Report CES-95-01). Department of Computer Engineering and Science, Case Western Reserve University, Cleveland, OH.

Yeap, Y. K. (1988). Towards a computational theory of cognitive maps. *Artificial Intelligence*, *34*, 297–360.